


OSG Blueprint Report & Storage Services

July 27, 2003

Ruth Pordes

The OSG Blueprint Activity




“This Blueprint for the Open Science Grid provides the guiding principals and roadmap for the building and operation of the infrastructure and will provide a basis for planning a coherent technical program of work.”

The Blueprint is not the actual plan.

An ongoing process to communicate to ourselves and others our end goals & architectural designs, exploring representative use cases and constrained by the technologies available now or possible in the future and the schedule of the participating groups.

An ongoing process to input to the planning and execution of the OSG.

Blueprint Activity Organization



Commitment of Contributors to meet “once a month for a week face to face”. To try and stimulate a core group who have sustained attention to the engineering and roadmap for OSG.


Broader circle of reviewers to try to ensure communication of blueprint can be successful and spread the understanding and collaboration.

Goal to distribute 1st draft of text to a broader community 15th August and present at the workshop in September.

Currently talking to people one on one based on output from the first face to face and having unstructured review phone meetings. Will meet with Wyatt for real time editing later this week.

New version sometime next week - don't want too many versions as people won't read it too many times. But it is clearly incomplete and incoherent to many readers :-). <http://listserv.fnal.gov/archives/osg-blueprint.html>.

OSG Blueprint, OSG-0 and Storage Services



Agreed that Storage Management is an essential component of OSG.

Agreed that Grid3 currently has no Storage Management Services.

Agreed that SE Grid interface is:


- SRM for the Management Interface;
- GridFTP for the file transfer service.

Agreed that Application (non-Grid) interface is “sub-set of Posix I/O”

US CMS Testbed has Storage Elements (SRM/dCache) at Tier-I and Tier-2s for use by US CMS VO and not yet available or supported for the multi-vo Grid3.

Agreed that OSG-0 required Storage Management at N sites. which are a mix of Strategic and Tactical, Persistent, Durable and Transient SEs

SE Issues:



Persistent, Strategic Mass Storage is supported by well provisioned Facilities. These are in hand at Fermilab, NERSC, BNL, JLAB - though with not enough effort available ?

Durable and Transient Storage Management is supported “spottily”

- SRM/dCache for the US CMS Sites Tier-1 and Tier-2 sites for the US CMS VO.
- LBNL DRM/HRM for STAR Tier-1/Tier-0 sites at BNL and NERSC.
- Parochial solution at UBuffalo using PVFS backend.
- In test at BNL - both SRM/dCache and LBNL/DRM

Grid3 includes & OSG-0 will include sites which support different VOs which may or may not present the same services. We will need to be clear on what the criteria are for deployment and support of basic OSG Storage Management Service at a Site - is it totally a Site responsibility or are there Common Infrastructure expectations and responsibilities also.

Support is needed to manage “batch queue local space for physical local file”; sandboxes with locally defined lifetime; and space reservation/use whose lifetime subject to policy.

Issues of file ownership on local site - what are the Grid user and VO expectations; what are the local application and site admin expectations; what are the responsibilities.

The Storage Element (SE) is responsible for saving/retrieving files to/from the local storage that can be a disk or a mass storage system. It manages disk space for files and maintains the cache for temporary files. In order to describe the proper service semantics, we need to introduce some concepts that will be applied.

Storage Elements represent a resource that may have very different quality of service metrics between different instances (see Figure 7.)




Figure 7: Space occupied by our different storage concepts in the space of QoS and Portability.

Opportunistic, 'tactical' Storage


At one extreme end we have Storage Elements that have a very simple purpose: store grid files close to a computing element as long as it's needed (if possible). Such SEs can be viewed as file caches, temporary or even scratch space. Their value is in their flexibility. Since they do not provide high data safety, such SEs can be easily deployed, switched on or off according to the need of a given site or a given virtual organization.

SEs are controlled by the sites and are subject to local policy. Having the concept of tactical SE allows a site to declare space in a local store in an opportunistic manner. It can provide storage that is currently unused by their local users and revoke it whenever necessary. This will make it attractive to sites to make resources available to grid users, knowing they can re-claim the resources whenever they want. Grid users profit from such storage being able to run jobs requiring local store at more sites. Users are expected to keep only disposable data in such stores - meaning that it should not matter to anyone if the instance of the data is lost - because it can be re-generated or re-copied from a master instance for example. Important data, master copies should not be kept in such storage (only at the user's own risk). If users generate new data at such stores, they should either register a master copy at a more long-term SE or be prepared to re-generate the data if necessary.

The tactical SEs may be hot-deployable and be alive only for a short period of time.

Fail-safe, 'strategic' Storage

Such storage comes with a higher quality of service. Users may expect to be able to reliably retrieve their files from such storage at any later time. Such store usually has a managed MSS behind it and virtual organizations are expected to pay the price of the storage, like the HEP experiments are expected to pay for Castor storage at CERN.



<http://opensciencegrid.org/techgroups/storage/index.html>

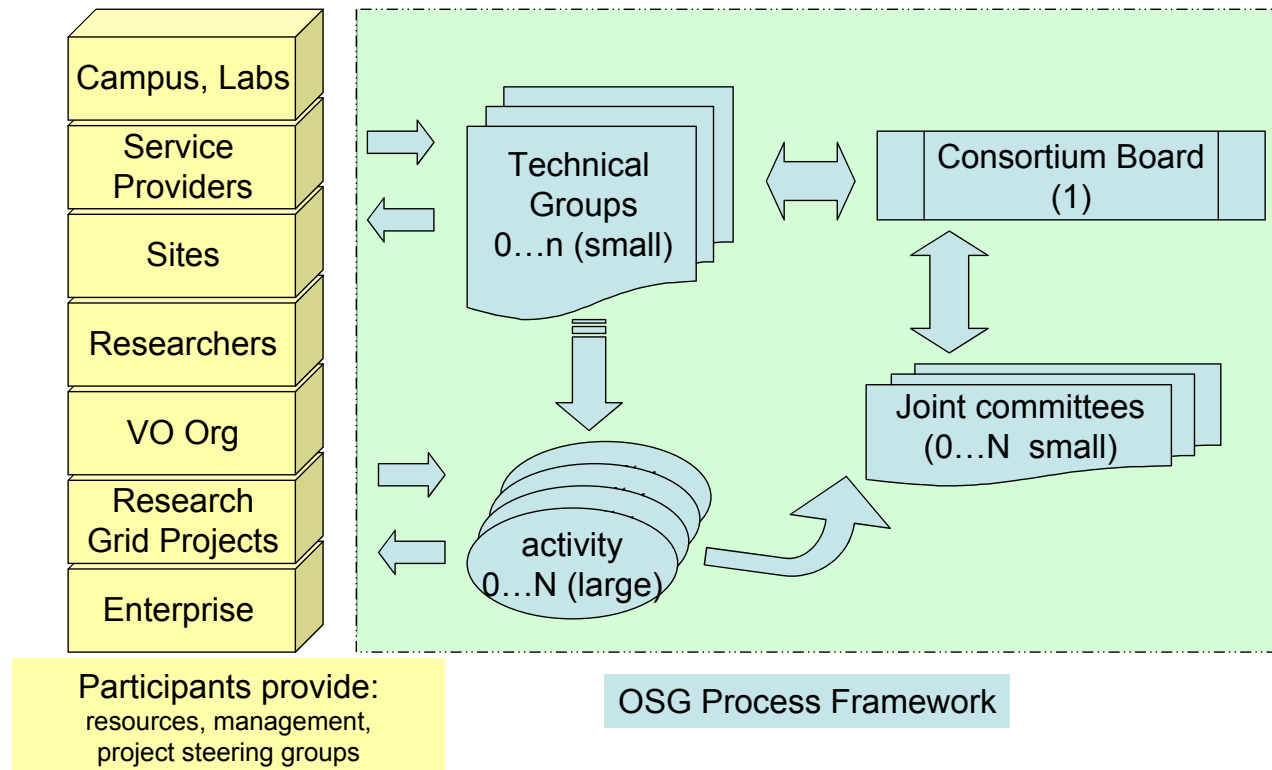

Storage Technical Group Mission

The Storage Technical group is responsible for coordinating the OSG activities that relate to data storage services. These services include:

- Provision of accessible storage resources of clearly-defined types (tape and disk), and flavor (permanent, persistent, transient, and volatile).
- Interfaces for the writing and reading of data by applications to and from storage, both over the grid and locally.
- Interoperability of storage interfaces across different administrative domains, initially OSG and EGEE.
- Published information about storage resources, including monitoring information, and agreements for their use.

U.S. ATLAS and U.S. CMS Tier-I centers will lead in the provision of their storage resources as part of a production quality grid, interoperating with other regional centers in their global experiment computing systems

Technical Groups within the OSG Framework




TG propose, organize and oversee Activities (=projects).

Liaise with their peer organizations in the U.S. and world-wide.

Participate in relevant standards organizations.

Representatives of each technical group participate in OSG-wide Joint Committees such as integration & operations.

Members of the Storage TG



gibbard@BNL.GOV	Bruce Gibbard
popescu@BNL.GOV	Razvan Popescu
shank@BU.EDU	Jim Shank
mlgreen@CCR.BUFFALO.EDU	Mark Green
miron@CS.WISC.EDU	Miron Livny
weber@CS.WISC.EDU	Jeff Weber
bakken@FNAL.GOV	Jon Bakken
bauerdick@FNAL.GOV	Lothar Bauerdick
ernst@FNAL.GOV	Michael Ernst
ifisk@FNAL.GOV	Ian Fisk
kennedy@FNAL.GOV	Rob Kennedy
petravick@FNAL.GOV	Don Petravick
rana@FNAL.GOV	Abhishek Rana
ruth@FNAL.GOV	Ruth Pordes
white@FNAL.GOV	Vicky White
skoranda@GRAVITY.PHYS.UWM.EDU	Scott Koranda
rwg@HEP.UCHICAGO.EDU	Rob Gardner
arie@LBL.GOV	Arie Shoshani
dlolson@LBL.GOV	Doug Olson
nlmeyer@LBL.GOV	Nancy Meyer
sychan@LBL.GOV	Steve Chan
huth@PHYSICS.HARVARD.EDU	John Huth
paul.sheldon@VANDERBILT.EDU	Paul Sheldon

Looking for a Chair / Coordinator to carry the program forward.

Current Activity is the Robust Data Transfer Service Challenge between CD/BNL/IT.

Need a full description of types of Storage Elements - still lots of discussion that needs to be recorded and checked as valid.